Opinions Sandbox: Turning Emotions on Topics into Actionable Analytics

Feras Al-Obeidat[1], Eleanna Kafeza[1], and Bruce Spencer[2](✉)

[1] Zayed University, Abu Dhabi, UAE
[2] University of New Brunswick, Fredericton, NB, Canada
bspencer@unb.ca

Abstract. The Opinions Sandbox is a running prototype that accesses comments collected from customers of a particular product or service, and calculates the overall sentiment toward that product or service. It performs topic extraction, displays the comments partitioned into topics, and presents a sentiment for each topic. This helps to quickly digest customers' opinions, particularly negative ones, and sort them by the concerns expressed by the customers. These topics are now considered issues to be addressed. The Opinions Sandbox does two things with this list of issues. First, it simulates the social network of the future, after rectifying each issue. Comments with positive sentiment regarding this rectified issues are synthesized, they are injected into the comment corpus, and the effect on overall sentiment is produced. Second, it helps the user create a plan for addressing the issues identified in the comments. It uses the quantitative improvement of sentiment, calculated by the simulation in the first part, and it uses user-supplied cost estimates of the effort required to rectify each issue. Sets of possible actions are enumerated and analysed showing both the costs and the benefits. By balancing these benefits against these costs, it recommends actions that optimize the cost/benefit tradeoff.

Keywords: Social commerce · Opinion extraction · Topic extraction · Actionable analytics

1 Introduction

Sentiment analysis and topic extraction have been very active research fields in recent years. The emergence of social media and the availability of a vast amount of user-produced information, make it possible to automatically identify users' emotions about the topics they discuss. Although data from a variety of social media have been analyzed for different purposes, the processing of data related to users' opinions as expressed in reviews has been of major interest.

Existing work has focused on providing useful information to users who want to check the sentiment expressed by others with respect to a business, service or product before buying it or marketers who want to know the general sentiment for their brands [5]. While the problem of topic extraction and discovering of sentiments polarity has been recently addressed in the literature [2–4]. there is

F. Belqasmi et al. (Eds.): AFRICATEK 2017, LNICST 206, pp. 110–119, 2018.
https://doi.org/10.1007/978-3-319-67837-5_11

no existing work to examine how topics and related sentiments can be used for providing actionable analytics for businesses. By actionable analytics, we mean data analysis and metrics that can enable businesses to better understand and improve their clients' opinions about their products and services.

In this work, we propose the *Opinions Sandbox*, a framework integrating topic extraction, sentiment analysis for extracting topics and their associated sentiments from an opinion database, and analytics for recommending actions. We propose a set of metrics that can contribute to developing exact strategies for influencing customers' opinions. Our methodology is based on existing popular techniques, including Latent Dirichlet Allocation (LDA) [6] for topic extraction, and the "bag-of-words" sentiment analysis algorithm where polarity is determined based on the frequency of occurrence of positive/negative words in a document. As a next post-processing step, we apply a procedure to incrementally inject into the corpus, comments that express opinions with the polarity of our choice. Thus we have a clear perceptive on how to influence the corpus. Although, business cannot inject opinions in social media, our Opinions Sandbox tool, gives insights on the number and style of opinions that could be generated for the specific topic and the specific customer base. Injected comments simulate the future situation after the business owner has taken steps to rectify the conditions that led customers to express negative opinions. This allows the business owner to determine the effect of various rectifying steps, and to decide which to do. We describe our running prototype that implements the these ideas.

We experiment on our ideas using the Opinosis 1.0 dataset [16] that contains reviews for many businesses. Our approach can be easily extended to any type of social media content. We use LDA and a frequency-based algorithm for opinion mining but our approach can work with any selection and combination of topic analysis and sentiment analysis algorithms. Our contribution is to provide actionable analytics for creating efficient strategies that can influence the opinions and the discussions of the customers.

The rest of this paper is organized as follows: we discuss related work, present the Opinions Sandbox idea, and describe an example implementation that provides recommendations that balance minimizing costs with maximizing social opinion. We discuss the main points and point to future work.

2 Related Work

While the problem of topic extraction has been addressed in the past, in recent years research is relating topic extraction with sentiment analysis. In their work [8] the authors are interested to mine users' opinions on Weblogs, analyzing the sentiments for subtopics. In their approach the authors propose a probabilistic mixture model called Topic Sentiment Mixture (TSM) where words are sampled by a mixture model of background language, topic language and two sentiment language models. They present a mechanism for extracting subtopics, associating with every subtopic a positive or a negative sentiment and how the opinions over a topic change over time. Their approach does not use LDA and the sentiment model is applied as a post-processing step to the topic discovery.

In their work [2] the authors present JST, a method that is using a weakly-supervised approach to draw words taking into consideration both topics and sentiment labels from a corpus of documents thus extending LDA. As a result, JST performs document level sentiment classification where topics and sentiments are detected simultaneously while it can extract sentiment oriented topics effectively evaluating the sentiment of each topic. In [1] the authors extend JST proposing the Sentiment LDA, where sentiment labels are associated to topics instead of documents and introduce sentiment dependency in their calculations. In more recent work, [11], the authors argue that the sentiment should not be used to influence the topic as done in JST but sentiment polarities as well as topics of text should be analyzed at the same time. They propose Double Latent Dirichlet Allocation (DLDA) for sentiment analysis in short texts. A review of on LDA-based topic extraction in sentiment analysis is presented in [7].

As already presented, there is a diversity of approaches in the literature regarding the extraction of topics and their associated sentiments. In our work, we are not attempting to provide yet another approach in addressing this problem, but we are making a next step and we are considering the following problem: given any approach on topic-sentiment analysis how can we improve the sentiment and how much will it cost to do so.

3 The Opinions Sandbox

Given a set of reviews about a specific product, service or business, such as a specific hotel in a city, we execute Latent Dirichlet Analysis [6], which partitions the reviews into disjoint sets. Each set pertains to a specific aspect mentioned in the reviews, such as the check-in experience, or the cleanliness of the rooms. LDA also generates a set of words associated with each set, from which it is possible to get some idea of the unifying themes and concerns discussed in a given set of comments.

Some of the reviews for a given topic express a positive sentiment, such as "The check-in procedure was a good experience." or "The cleanliness of the room was very satisfactory", while others express a negative sentiment. Some reviews express a combination of views. There are a variety of techniques to assess the overall sentiment or mixture of sentiments attached to a review, and also to any set of reviews, including the entire set. The proposal in this paper is a framework that depends on the existence of some method of extracting the sentiment, but is not dependent on any particular method.

The Opinions Sandbox uses the subsets of reviews as generated by LDA partitioning, and the sentiment assessed for each partition subset, to generate actionable recommendations aimed at addressing the issues mentioned in the reviews and improving future sentiment analysis results. The Opinions Sandbox is a recommendation system to be used by business owners to address their clients' concerns and thereby improve their online profile.

The business owner works cooperatively with the Opinions Sandbox to identify the issues that are both addressable and damaging to the overall sentiment.

After the reviews are partitioned by LDA, the sentiment of each partition is assessed. The business owner can see which partition is least positively assessed and can see the unfavourable reviews in that partition, starting with the least positive. He or she assesses the issues mentioned in the partition assesses whether and to what extent each can be addressed. The Opinions Sandbox generates positive reviews in sufficient number and strength to counter the negative reviews that the business owner is considering to address. Given the generated positive comments, the sentiment within the Opinions Sandbox will rise for the topic. The business owner continues until sufficient positive reviews are generated so that sentiment assessment achieves a level that is satisfactory to the business owner. This gives the business owner a clear understanding of the degree of work that needs to be done in order to redress issues mentioned by customers in the reviews.

Thus the Opinions Sandbox provides an aid by which the business owners can efficiently traverse the set of online reviews and quickly identify any deficiencies in their business offering. First, it use partitioning to cluster the issues on similar topics. Second it assessing the sentiment on each topics. Third it takes advice from the business owner who estimates the degree of work needed to address those issues. Fourth it predicts the overall change in the online opinion after that work is done, by creating simulated positive comments that might arise when the work is done, and re-assesing the sentiment that would result. Finally, it presents a set of options ordered by degree of work, which the business owner can consider as a set of recommendations that maximize his or her return on investment while addressing issues that aggravate customers and negatively affecting the online reviews.

4 An Opinions Sandbox Example Implementation

The Opinions Sandbox is a framework and we also provide a specific implementation in R/Shiny [12]. The framework depends on an LDA-based partitioning of comments and on some technology for assessing sentiment. The implementation uses a concrete method for partitioning documents in the `tm` package [13] and the `topicmodels` package [14]. Topic selection is done both by Gibbs [10] and CTM [9] techniques. The user is shown both results. While there are a variety of techniques for assessing comments, the implementation currently uses a straightforward opinion assessment based on assigning sentiment to certain words, either positive and negative. The sentiment of any collection of words, whether a single comment, a cluster of comments, or the entire set of reviews, is based on the number of occurrences of positive and of negative words. The Opinions Sandbox implementation is illustrated using comments collected online and made available publicly from the Opinosis 1.0 Dataset [16]. This dataset consists of 51 files, each containing about 100 comments, selected from various sources.

More specifically, our opinion assessment strategy applies to a single review or to a set of reviews. It considers each review or set of reviews to be a bag of words, which is a multiset, *i.e.* a set of word instances. We use a given set P of words associated with positive sentiment such as "good" and "satisfactory", and a set N of negative words. These categories are provided by the General Inquirer dataset [15]. Given a bag B of word instances, comprising either a single review or a set of reviews, we assess each word instance as either having no sentiment or of having positive or negative sentiment according to whether it belongs to the positive set P or to the negative set N. We associate a sentiment metric S based on each word occurrence: S is the fraction of sentiment words in B that are positive.

$$S(B) = |\{B \cap P\}|/|\{B \cap (P \cup N)\}|$$

Similarly we compute the $N(B)$ to measure the negativeness of the corpus B.

$$N(B) = |\{B \cap N\}|/|\{B \cap (P \cup N)\}|$$

In the Opinions Sandbox framework, it is suggested to simulate the effect of addressing client's issues by creating and injecting positive comments into the reviews. However, in our initial implementation we circumvent this step, because of the simplicity of the sentiment assessment. When the business owner identifies a set of comments to have been addressed, we simulate the result of that work having been done by increasing the positive word count by a number equal to the negative word count. In effect, this simulates having each negative opinion countered by a new positive opinion. For instance, suppose there are 100 sentiment words in a set of reviews, of which 40 are positive and 60 are negative. Given our assessment method, this would generate an assessment of $40/(40 + 60) = 0.4$. Once the business owner deems them as addressable, the future assessment is predicted to be $(40 + 60)/(40 + 60 + 60) = 0.625$, as if 60 new opinions, each expressing a positive sentiment, were added. This strategy will always increase the sentiment of a given set of comments, and will always convert an assessment below 0.5 to one above 0.5.

We also consider a more powerful injection mechanism, where the positive comment completely counters the effect the negative comment. In our previous example with 40 postive and 60 negative comments, the effect of adding 60 new postive comments to counter the negative comments brought the sentiment to 0.625. However, if the positive comments nullify the negative comments, the new sentiment is $(40+60)/(40+60)$. Currently our system uses this mechanism. The goal, then, is to create a sentiment of 1.0 for the whole corpus.

5 System Description

The flow of the Opinions Sandbox system is described in Fig. 1. The user is free at any time to restart at any section, for instance, to consider a new business, or analyse a new partition. Screen shots of the running system are shown in Figs. 2, 3, 4, 5 and Table 1, which shows the user accessing various parts of the flow mentioned in Fig. 1.

We consider a specific example of a hotel in the San Francisco area. During the review, the business owner is deemed to want to address the comments in

1. Select the business, service or product, Figure 2.
2. Load the comments.
3. Select the number of topics of interest, or use the suggested topics
4. Partition the comments into topics, Figure 3.
5. For each partition:
   a) Review the comments, Figures 5.
   b) Estimate and provide the cost of addressing all comments in each topic.
6. Review the report telling the effect of addressing comments, the cost of doing so and the effect on sentiment, Table 1.

Fig. 1. User's steps: flowchart for opinions sandbox

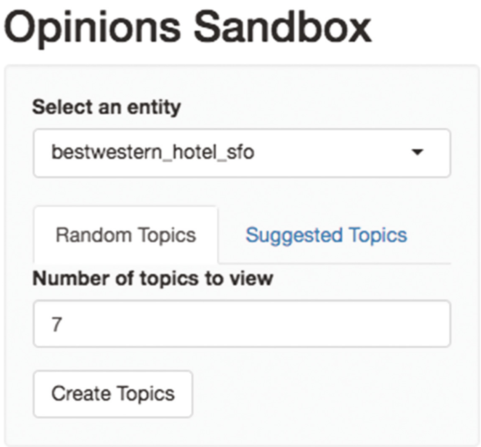

Fig. 2. The start screen allows the user to select reviews for a business, product or service and to choose a number of topics. Thus user can also select from a precomputed selection of topics from this dataset, which is available within the Opinosis 1.0 dataset.

each of the seven topics. The cost of doing so is estimated to be 500, 400, 800, 800, 1200, 1400, and 800, respectively, in some unspecified monetary units. We do so by placing different weight on each of the two criteria, and by ordering the various combinations either by sentiment or by cost. Table 1 show the partial enumeration of the 2^n possible choices of addressing or not addressing each of the $n = 7$ topics. If nothing is done, the cost is zero and the sentiment is predicted to remain at 0.73. If all of the comments are addressed, the cost is 5900, but the online sentiment will have a positive comment for every negative comment, and thus under the stronger form of injection, achieves a high score of 1.

We also provide in this paper the cost and expected resulting sentiment for each of the 2^7 combinations of the the seven sets of comments, in Table 1. As that

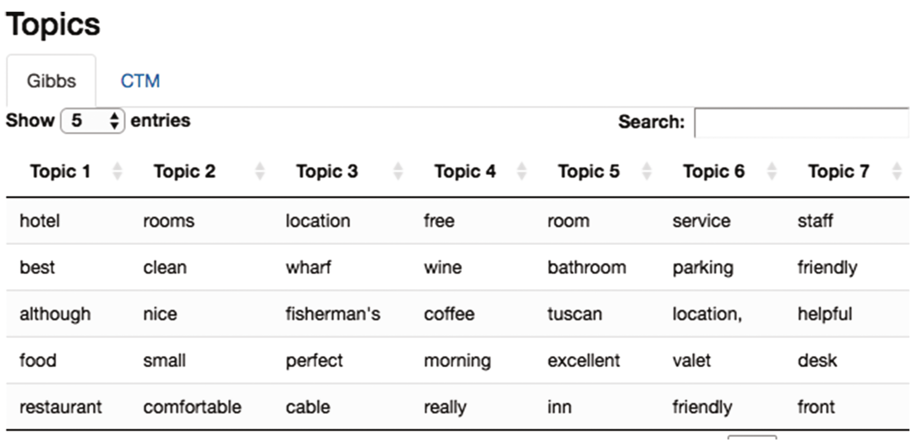

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7
hotel	rooms	location	free	room	service	staff
best	clean	wharf	wine	bathroom	parking	friendly
although	nice	fisherman's	coffee	tuscan	location,	helpful
food	small	perfect	morning	excellent	valet	desk
restaurant	comfortable	cable	really	inn	friendly	front

Fig. 3. Once the number of topics is selected, the user can review the topics that were selected. The user can choose either Gibbs or CTM sampling for partitioning into topics.

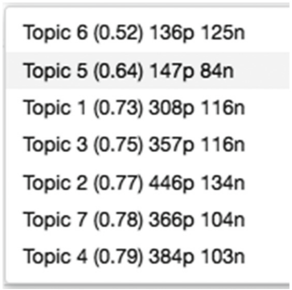

Fig. 4. Within a drop down selection list, the user can see the current sentiment for each topic, as well as the number of positive and negative comments. The topics with lowest sentiment are at the top of the list.

table shows, the result of addressing all of the comments is an absolutely positive online sentiment, assuming the stronger form of comment injection where the new positive comment is assumed to override the existing negative comment.

The system then blends two criteria in Table 1, cost and benefit, to make a recommendation of which jobs to do. The blend can be oriented toward lower cost by slightly weighting the cost criterion. In this case the system recommends addressing the comments for topics 1, 2 and 3 at a cost of 1700 to raise the sentiment to 0.86. The system can also be tuned to consider higher sentiment as more important. In this case, the system recommends addressing comments in topics 1, 2, 3, 4 and 7 at a cost of 3,300 and raise the sentiment to 0.93.

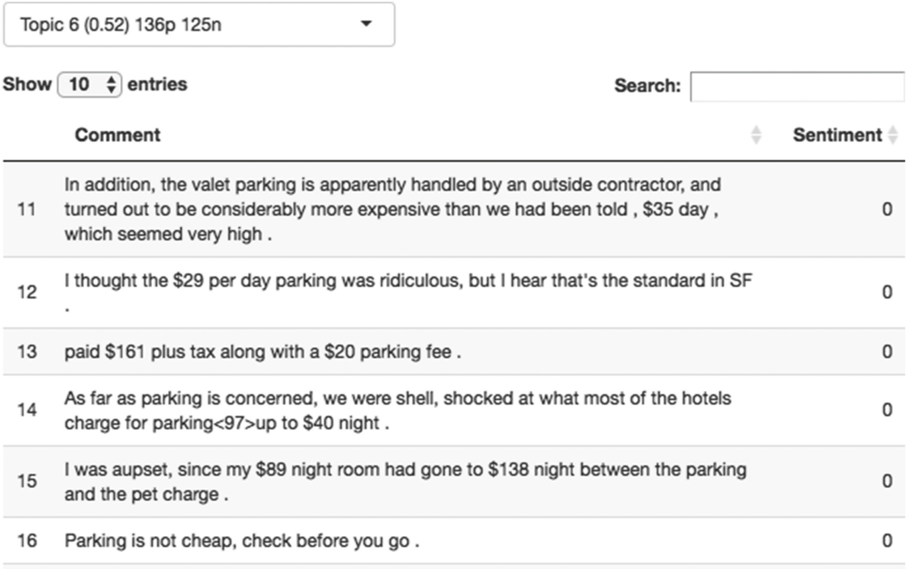

Fig. 5. Comments are presented along with their sentiment. Parking issues elicit negative sentiment in Topic 6.

Table 1. Description and cost for each job, and the resulting sentiment

Job combination	Cost	Sentiment result
do nothing	0	0.73
t2	400	0.78
t1	500	0.77
t3	800	0.77
t4	800	0.77
t7	800	0.77
t1 + t2	900	0.82
t5	1200	0.76
t2 + t3	1200	0.82
t2 + t4	1200	0.81
t2 + t7	1200	0.81
t1 + t3	1300	0.81
t1 + t4	1300	0.81
t1 + t7	1300	0.81
t6	1400	0.78

(*continued*)

Table 1. (*continued*)

Job combination	Cost	Sentiment result
t2 + t5	1600	0.81
t3 + t4	1600	0.81
t3 + t7	1600	0.81
t4 + t7	1600	0.80
t1 + t5	1700	0.80
t1 + t2 + t3	1700	0.86
⋮	⋮	⋮
t1 + t2 + t3 + t4 + t7	3300	0.93
⋮	⋮	⋮
t1 + t2 + t3 + t5 + t6 + t7	5100	0.96
t1 + t2 + t4 + t5 + t6 + t7	5100	0.96
t2 + t3 + t4 + t5 + t6 + t7	5400	0.96
t1 + t3 + t4 + t5 + t6 + t7	5500	0.95
t1 + t2 + t3 + t4 + t5 + t6 + t7	5900	1.00

6 Conclusions and Future Work

The Opinions Sandbox is a tool that helps business owners to assess the severity of online criticisms. It first partitions the set of online reviews according to topic. Each partition pertains to one or a small set of issues to which the business owner can respond. After working with the business owner to identify the issues, how to resolve them and the degree of effort required, the Opinions Sandbox then injects positive comments that counter the effect on the existing negative comments, thus simulating future situation where the issues are addressed. It allows the business owner to contrast the current online sentiment with a forecast of the future sentiment. The Opinion Sandbox helps the business owner to quantify the amount of work to address issues mentioned, and the result that doing so is likely to have on the online opinion. It enumerates the combinations of actions that can be taken, and the effect of each on the online opinion, so that the most cost-effective method can be found for addressing some or all of the issues. In summary, the Opinions Sandbox helps the business owner to quickly understand the online issues, to consider the possible redress actions, and to find a selection of actions that provides the most expedient way to improve online sentiment.

This product is particularly relevant for developing economies, and in regions including the Middle East and North Africa, where online tourism attracts potential customers making their first visit. These customers rely heavily on online recommendations, sometimes only on these recommendations, before making significant purchases. In future we will experiment with different topic classification techniques, and with different techniques for measuring sentiment. Comment synthesis is a relatively new area and we plan to contribute. We are also planning trials with clients in the tourism industry, where opinions have direct economic impact.

References

1. Li, F., Huang, M., Zhu, X.: Sentiment analysis with global topics and local dependency. In: Proceedings of AAAI, pp. 1371–1376 (2010)
2. Lin, C., He, Y., Everson, R., Ruger, S.: Weakly supervised joint sentiment-topic detection from text. IEEE Trans. Knowl. Data Eng. **24**(6), 1134–1145 (2012)
3. Chen, X., Tang, W., Xu, H., Hu, X.: Double Lda: a sentiment analysis model based on topic model. Paper presented at the 2014 10th International Conference on Semantics, Knowledge and Grids, 27–29 Aug 2014
4. Yin, S., Han, J., Huang, Y., Kumar, K.: Dependency-topic-affects-sentiment-Lda model for sentiment analysis. Paper presented at the 2014 IEEE 26th International Conference on Tools with Artificial Intelligence, 10–12 Nov 2014
5. Go, A., Bhayani, R., Huang, L.: Twitter Sentiment Classification Using Distant Supervision. Technical Report, Stanford University (2009)
6. Blei, D.M., Andrew, Y.N., Michael, I.J., Lafferty, J.: Latent Dirichlet allocation. J. Mach. Learn. Res. **3**(4/5), 993–1022 (2003)
7. Rana, T.A., Cheah, Y.-N., Letchmunan, S.: Topic modeling in sentiment analysis: a systematic review. J. ICT Res. Appl. **10**(1), 76–93 (2016)
8. Mei, Q., Ling, X., Wondra, M., Su, H., Zhai, C.: Topic-Sentiment Mixture: Modelling Facets and Opinions in Weblogs, pp. 171–180 (2007)
9. Blei, D.M., Lafferty, J.D.: A correlated topic model of science. Ann. Appl. Stat. **1**(1), 17–35 (2007)
10. Phan, X.H., Nguyen, L.M., Horiguchi, S.: Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In: Proceedings of the 17th International World Wide Web Conference (WWW 2008), Beijing, China, pp. 91–100 (2008)
11. Xue, C., Tang, W., Xu, H., Hu, X.: Double Lda: a sentiment analysis model based on topic model. In: Proceedings of the 2014 10th International Conference on Semantics, Knowledge and Grids, pp. 49–56. IEEE Computer Society (2014)
12. http://shiny.rstudio.com/
13. Feinerer, I., Hornik, K.: https://cran.r-project.org/web/packages/tm/tm.pdf
14. Grün, B., Hornik, K.: https://cran.rproject.org/web/packages/topicmodels/topicmodels.pdf
15. http://www.wjh.harvard.edu/inquirer/
16. Ganesan, K.A., Zhai, C.X., Han, J.: Opinosis: a graph based approach to abstractive summarization of highly redundant opinions. In: Proceedings of the 23rd International Conference on Computational Linguistics (COLING '10) (2010)